Trade-offs involved in the choice of Cloud Service Configurations when Building Secure, Scalable, and Efficient IoT Networks

Amitabh Mishra, Department of Computer Science, University of West Florida, Pensacola FL 32514 USA
Thomas Reichherzer, Department of Computer Science, University of West Florida, Pensacola FL 32514 USA
Ezhil Kalaimannan, Department of Computer Science, University of West Florida, Pensacola FL 32514 USA
Norman Wilde, Department of Computer Science, University of West Florida, Pensacola FL 32514 USA
Ruben Ramirez, Department of Computer Science, University of West Florida, Pensacola FL 32514 USA

Corresponding Author: Amitabh Mishra, Ph.D.
Address for correspondence: University of West Florida, 11000 University Parkway, Building 4/232, Pensacola FL USA 32514
Email address: amishra @uwf.edu
Telephone Number: +1-850-473-7346

Total word count: 5850 words

Abstract length: 192 words

Funding: This work has not been funded by any agency.

Conflicts of interest: The authors do not have any conflicts of interest to report.

***Abstract*— This article focuses on results obtained from two cloud-based models that examine trade-offs between security, scalability and efficiency of data collection for IoT sensor networks. This work can provide insight for IoT systems designers in choosing security controls and scalability features when working with cloud services. The results were obtained from a smart home IoT prototype system**

in which data records from in-home sensors are transmitted wirelessly to an in-home hub, which forwards them to a Cloud Web Service for storage and analysis. We consider different configurations and security controls on the wireless (in-home) and on the wired (home-to-web) sides. The configuration on the wireless side includes encrypted or plain-text transmission from the wireless sensors to the in-home hub for probing if software encryption of sensor data adds appreciable delay to the transmission time. The configuration on the wired side includes encryption or plain-text transmission, with or without authentication, with or without scalable Cloud services. For each configuration we measure end-to-end latency, transmission latency, and processing latency at the web service. Results of the experiments on the wired side showed much greater latencies and variability of latencies when using scalable Cloud services.

Index Terms— **Internet of Things, Cloud Computing, Security, Scalable Services, Encryption**

I. Introduction to Smart Homes and the Internet-of-Things

SMART Devices and the Internet-of-Things (IoT) are increasingly pervading our daily lives shaping and transforming our experiences with the physical world. In the home, smart devices add conveniences, comfort, or safety. They regulate environmental conditions in the home by tracking our activities to infer our needs for cooling and heating while optimizing energy consumption. They keep track of our fitness & health conditions [1] or provide new services for controlling and maintaining devices in the home [2]. Market studies forecast an exponential growth of IoT consumer devices with more than 13 billion new devices connected to our networks by 2020 [3]. Those devices may leverage wireless networks available in a home to connect to the Internet and integrate with Cloud services. This integration enables smart devices to access software and firmware updates as they become available, tap into vast Cloud computing and data storage services, or offer automation and remote monitoring and control for residents of a home.

For smart homes to adapt to user needs and provide conveniences, comfort, and safety in the home, data from the smart devices must be shared and integrated to build a complex model of a home and its residents. These models can be used to track activities and assess the resident's needs and make recommendations or take actions in situations where an emergency arises. For example, a smart house that detects a water leak in the basement while the resident is at work can instantly notify the resident via text messages while also shutting off the main water valve of the house to prevent flooding in the home.

With limited data storage and processing capabilities, smart devices normally rely on Cloud services to handle their data. Those services provide the scalability and reliability to ensure that the Smart Home is responsive to the resident's needs. Smart devices either directly connect to Cloud services to upload their data or use a hub in the home that collects the data from multiple Smart Home devices and uploads them to the Cloud (see Figure 1). As devices vary in their needs to communicate with Cloud services to send data back and forth, we can expect a wide range of communication requirements for bandwidth and latency.

Studies have shown that smart IoT devices offer minimal to no protection from cyberattacks that may compromise the device or the data that it stores or transfers into the Cloud [4]. Security related standards and regulations are not well established yet, resulting in numerous legacy communication protocols which may not provide ample protection to data privacy. Moreover, the overhead of such legacy protection for data storage and transmission is not well understood.

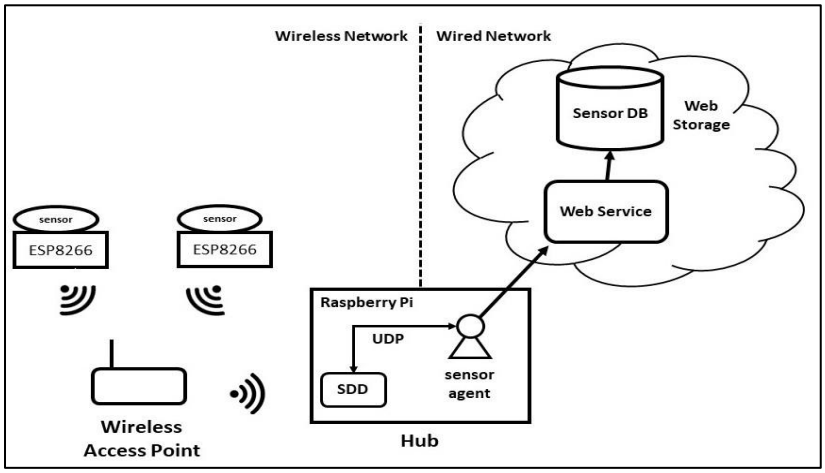

Fig. 1. IoT / SMILE Home Network Overview.

In response to these challenges, as a continuation of our previous work in [5], we report on two cloud-based models that examine design alternatives that provide trade-offs between security, scalability and efficiency for IoT networks. The study was performed on a prototype Smart Home system using multiple sensors in a network developed at the University of West Florida. This sensor network is also utilized in another research project intended at developing a Smart Independent Living for Elders (SMILE) home

[6].

We used this experimental platform for our two cloud-based models. The first considered different security controls for the wireless part of the network on the left side of Figure 1. The second considered different controls and designs for the wired part on the right side of the figure. In both cases, for simplicity of measurements, we consider unidirectional data flow from the sensor, via the Sensor Agent and the Hub, to the Web Store which consists of a set of several web services with a backend data storage. As our indicator of efficiency, we use latency, defined as the overall elapsed time between message generation and receipt. We used the Network Time Protocol (NTP), to guarantee synchronization among clocks between the different system components.

In this article, section II provides background on the cloud-based model, specifically the security issues that are being studied and those not studied. Section three describes the architecture of the prototype SMILE home system including the wired and the wireless components in the system and the web services and data storage facilities on which it relies. Section four describes related work on IoT security while section five and six describe the cloud-based model for the wireless and wired section of the system. Finally, section seven discusses the experimental results and conclusions of the cloud-based model.

II. Smart Home/IoT Security and Scalability Context

The adoption of IoT networks in general, and of Smart Home networks, introduces significant security risks. In this paper we consider design alternatives that address an important subset of these risks. In the terminology used by NIST [7], we consider an outsider adversary with high to very-high capabilities, such as an experienced and well-financed hacker group. Such an adversary may be seeking to capture valuable data about Smart Home occupants or to compromise network components for use in denial of service attacks. The adversary will have the capability to intercept and interfere with wireless and possibly wired communications (see Figure 1). However, we exclude from consideration insider attacks on the producers of IoT components since they are not addressable by network design but rather by other means such as a careful selection and monitoring of personnel. Similarly, we will not consider physical attacks on the sensors or other devices within the Smart Home since these are best addressed by physical security and by introducing tamper-proof components into each device. Finally, we largely exclude attacks on vulnerabilities in the web service itself or in the web storage facility since these classes of attack are not specific to IoT applications. Thus, our focus is on evaluating design approaches that provide secure communications, both wireless and wired within a home.

We organize our analysis of the different design approaches around the traditional CIA Information Security Triad of confidentiality, integrity, and availability attributes [8]. Related to availability is a fourth attribute of scalability; in many cases IoT networks will need to sustain availability as data volume grows, either gradually as more nodes are added or at extreme events such as a flood or a strike.

Confidentiality is an essential attribute in almost all IoT implementations. In a Smart Home for example, sensors measuring weight, illumination or temperature can divulge living habits of the occupants, effecting their privacy. Availability and integrity become prominent considerations contingent on the use case. Failures in monitoring the refrigerator to reorder milk would be more a nuisance than a disaster. However, if a fall detector on a resident in an assisted living home is programmed to call for an emergency ambulance service, system availability must permit continued operation even in DoS attack scenarios while data integrity ensures defense against spoofing.

The design approaches we used also impact the scalability of the IoT network, but the scalability obligations depend significantly on the IoT application. If the smart homes are located within a single gated community, the associated volume of data can be anticipated. On another hand, Smart-city-wide IoT applications must be able to handle huge and flexible data volumes with great reliability. Such challenges can be better addressed through cloud infrastructures.

III. The SMILE Home System Architecture

The design alternatives we evaluate in this article were implemented within the general IoT network architecture called Smart Independent Living for Elders (SMILE) that is shown in overview in Figure 1. An important design criterion for this architecture is the ability to update after deployment. A Smart Home must be able to incorporate new sensors and technology without disruption. Changes to the sensor hardware that enhance system performance and reliability must disrupt the system's response, data collection, or processing. Similarly, for the software, bug or security fixes or needed upgrades must be installable without disruption.

Another important design criterion is that significant computation capacity may be needed to correlate sensor data and make any needed near-real-time decisions. However, the devices in the home are unlikely to have appreciable storage and processing power, so most processing must be moved to the Cloud.

To meet these goals, the SMILE architecture is created using distributed processes modeled through software agents which either run as part of the web service in the Cloud or execute in the home on the home's central computing device, which acts as a hub to collect sensor data and or interact with residents in the home. All SMILE devices like sensors, widgets for browser visualization, or actuators are operated using software agents. Each agent wraps the technology and offers an interface for data collection and device control to generate a Smart Home response. To handle updates, new or updated agents would be signed to establish authenticity and pushed to the hub or the Cloud.

The sensor boxes in our experiments are created using commonly available sensors and ICs such as single board computer nodes, temperature, light, pressure, and infrared sensors. Separate sensor boxes are assembled using electronic components for creating the sensors. The sensors gauge environmental data and the ESP8266 IoT node in each box collects and sends the data using the MQTT (Message Queuing Telemetry Transport) application layer protocol for reliable data delivery. All wireless devices transport data in the home via a wireless access point to the Raspberry Pi, which acts as the sensor hub.

The data packets sent by the mesh network are received by a Raspberry Pi platform running the Sensor Data Distribution (SDD) software on a continuous basis. The packets are then forwarded by the SDD over an internal UDP channel to Sensor Agents which are sensor processing units running on the Pi platform. The Sensor Agents can be restarted remotely, replaced or shut down as required in the process of Smart Home data collection.

The Sensor Agents running on the Pi hub have wired Internet communication to a web service in the Cloud. For the experiments reported on in this paper, the web service simply stored the sensor values in a Sensor Database. In a complete SMILE system, the web service would process the accumulated Smart Home data utilizing machine-learning algorithms to deduce activities in the home in order to provide smart assistance for assisted living as required. The web service and the Sensor Agents are implemented using Java EE and Java respectively.

IV. Related Work on IoT Network Security

Ubiquitous services in healthcare, logistics and unmanned systems are being made possible by connecting real-world things to the Internet for communications in form of IoT devices that are growing in numbers every day. Bain and Co. forecasts that the Internet of Things market would grow to $520 Billion in 2021 [3]. IDC estimates the spending on IoT to reach 1.2 Trillion by 2022. Modular, application specific IoT devices in market today incorporate dedicated hardware, network connectivity to the cloud for data collection and processing.

Before the general public starts trusting IoT products, multiple technological and social hurdles need to be crossed [9]. The prominent problems for IoT acceptability that stand out include making them smarter by facilitating their adjustment and autonomous performance while ensuring security, privacy and trust, and empowerment of total interoperability of interconnected IoT gadgets. A basic consideration for IoT is the security of personal data which Is gaining prominence with the growing awareness amongst the masses on loss of privacy [10]. While some possible standards have started to emerge on the horizon, now, there still seems to be no consensus amongst IoT product manufacturers regarding the implementation of security in IoT gadgets [11]. The US National Intelligence Council lists IoT as one of the six civil technologies that are potentially disruptive and could impact US national power [12]. The worst mentionable security lapses happened when millions of healthcare records were leaked [13], the signaling system for the CSX Train network was attacked [14], a variety of new vehicle models from prominent vehicle manufacturers were hacked into, forcing a recall hacking of a car leading to recall [15], altering the readings of smart energy meters [16] and bringing down of a major portion of the Internet through DDoS attacks from IoT baby monitors [17], amongst several others. Such incidents imply that serious security considerations need to be inserted in the process of Smart Home design. The variety of products and data on the cloud platform are devoid of restrictions on fixed infrastructure and security owing to cloud computing features such as location transparency, service abstraction, and dynamic scalability [18]. Security considerations become crucially important on the cloud owing to enormous storage requirements combined with fast access and information processing [18].

A detailed survey on the considerations and issues in IoT security is offered by Jing et. al [19]. The article also investigates and suggests ways to heterogeneously incorporate IoT products and services across layers. Security becomes a vital design objective for data-in-transit security, where data are encrypted and hashed before transmission over the Internet. This has a considerable influence on performance which has been overlooked by researchers so far [20]. Müller et al describe a work to ascertain the performance impact of security strategy decisions in arbitrary cloud database systems and of using certain combinations of encryption and hashing algorithms in Cassandra and DynamoDB [21].

Motivated from the innovations offered by cloud computing, Kardas et. al proposed a security and privacy model for RFID systems merged with the cloud computing model to enhance the scalability and performance while preserving the security and privacy of the systems [22]. The work reported by Sun et al primarily aims to emphasize the major security, privacy and trust concerns in current cloud computing systems and assist users realize the tangible and intangible threats linked with using them [23].

Few works in the literature studied the various tradeoffs in IoT environments. The authors of [24] use an unreliable link model to investigate the tradeoff between the energy consumption and wireless communication in a multihop network. The framework in [25] surveys Delay-Tolerant Networks (DTNs) for tradeoffs between the packet delivery delay and types of packet transportation costs in wired network and mobile wireless network environments.

The research presented in [26] discusses the wide-stretching application security demands for IoT and varied network-edge resource limitations for end-to-end cloud-fog communication. The authors present a secure end-to-end protocol for resource-limited devices by adjusting security functionality used in unconstrained devices. Similar work was reported in [27], which analyzes the security problems of each layer in the TCP/IP model and the issues related to cross-layer heterogeneous integration and security.

V. Wireless Network Experiment

The instrumentation setup starts with the wireless sensors that sense ambient physical parameters and relay the data to the next stage in the sequence, the sensor agent over a wireless link. The following section explains the details of the wireless network configurations.

A. Experimental Setup

A Smart Home involves a limited, manageable number of IoT sensors and actuators making scalability in the wireless network a non-issue but the same cannot be asserted about the integrity and confidentiality aspects of the application data. The IoT network is prone to spoofing and sniffer attacks for collecting unauthorized data and disruption of applications and services. Resorting to encryption of IoT sensor and actuator data using a shared key working the sensor and the Sensor Agent could take care of integrity and confidentiality.

Sensor Agents are deployed to receive transmitted sensor data relating to physical quantities on wireless links from sensor nodes in a Smart Home. By default, these transmissions are plaintext and are vulnerable to tapping by a tuned receiver, making listening and man-in-the-middle attacks a potential threat. To add application level security and reduce the risks of such threats to this data we encrypted it prior to transmission. Due to high frequency and hence high amount of generated data, the time taken at the smart sensor nodes to encrypt sensor data could add additional encoding latency in addition to local processing delay. In our experiment, we measured the mentioned delay in communication post encryption and contrasted the throughputs with and without encryption.

Managing the energy consumption in wireless IoT sensors and actuators is an important consideration for Smart Home designers because of battery based power of such sensors and actuators. The limitation of powering remote IoT sensors and actuators makes it essential for the designers to also keep battery life and the number of replacements in perspective. The power consumption in sensors involves uses in transduction, analog to digital conversion, further processing followed by transmission and reception. Encrypting sensor data would mean more power required for processing, adding to the power consumption. This additional energy consumption due to encryption can be found out by checking the battery drain. This battery consumption can be optimized [28].

In the experiments for the wireless network configuration, we used ESP8266-based sensor boards. The ESP8266 is a low-cost 32-bit Wi-Fi enabled microcontroller chip running full, integrated TCP/IP protocol stack. It requires nominal external circuitry and incorporates the microcontroller, standard digital peripheral interfaces, filters, amplifiers, RF balancing unit, antenna switches and power management units in the same package. It can be used to provide cost-effective, low-power, highly-integrated Wi-Fi solutions.

The ESP8266 System-on-Chip (SoC) microcontroller board comes with a 2.4 GHz Wi-Fi (802.11 b/g/n, supporting WPA/WPA2) built into it. It is apt for building applications for wireless sensor networks and WiFi enabled IoT systems. Our experimental setup used a specific model of ESP8266, the ESP-12E to study the performance of an IoT node in encrypting and sending a stream of data over the TCP/IP protocol stack.

Our IoT network utilizes the Message Queuing Telemetry Transport (MQTT) protocol as the application layer protocol for reliable data delivery [29]. MQTT is a widely used protocol best suited for message delivery to remote locations that have low network bandwidth and/or low processing capabilities. It is a Publish-Subscribe protocol which means that any nodes subscribed to a topic on the MQTT broker receive all the messages published to that topic by any other node that publishes to it. This project uses the Mosquitto as the MQTT broker and PubSubClient as the client.

The network was setup in two different configurations to measure the latency, reliability and consistency of the data transmitted over the network. Each of these two networks consist of a sensor node running the pubsubclient, a WiFi router and a receiving node that hosts the MQTT broker. The two configurations are as follows.

1. ESP8266 (Sensor Node) – Netgear N450 (WiFi router) – Raspberry Pi 3 (Receiving Node): The Raspberry Pi 3 was installed with Mosquitto broker and configured to listen at 1883 for MQTT connections, as in Figure 2.
2. ESP8266 (Sensor Node) – Android Smartphone Mobile Hotspot (WiFi Router) – Raspberry Pi (Receiving Node): An Android smartphone was configured to serve as a WiFi hotspot at a unique SSID to receive connections from only known hosts. This setup was used to simulate an isolated network for getting unaffected readings for the latency of the data transfer, as shown in Figure 3.

In both the network configurations, the ESP8266 was flashed with the same code involving the transfer of a dataset comprising of 10000, 24-byte string values. In each case, after establishing a WiFi connection with the router, ESP-12E would start sending the assigned number of 24-byte strings over MQTT after one or two rounds of AES CBC encryption on every string of data. An

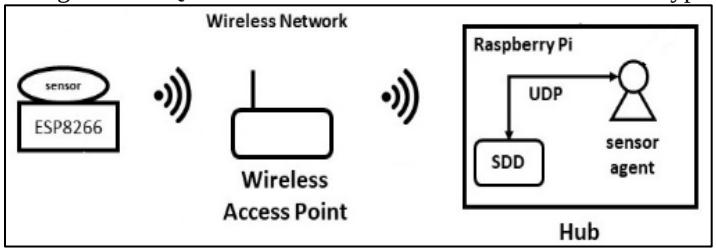

Fig. 2. IoT Wireless Sensor Network Configuration 1

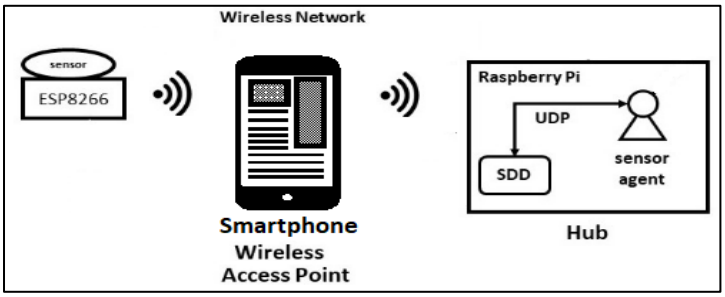

Fig. 3. IoT Wireless Sensor Network Configuration 2

AES library was used for encrypting the 24-byte strings. For both the configurations, the same size datasets were used to obtain the latency values for the data transfer. Multiple iterations were performed on both the configurations to obtain statistical consistency in the readings. Several observations were made regarding the behavior of the IoT network depending on the network activities of the home-based network, considering it a traditional smart-home system network.

B. Experiment Results

Like several other off-the-shelf sensors commonly available, the sensors that we used for building the IoT Smart Home did not come with in-built hardware security. Security was an add-on feature to these sensors. To get a base time for raw, unencrypted data was sent and the transmission times were noted for the sensor values across 10000 messages to get a good average. This was done for two different media using a home router for connection in one setup and the cellphone hotspot in the other. Five different runs of the experiment were performed to ascertain a range on variability of transmission time. Transmission times were also noted for the same number of messages in five different runs for encrypting and then sending the data. The data was then encrypted a second time to improve on the security and times were noted as before. Table I and Figure 4 show the findings of the experiments for the two connectivity media arrangements. As expected, a delay on the transmission time was found in case of single encryption, which increased for double encryption for both the media. However, it was observed that the overhead for encryption through software wasn't too high and was negligible for most real-time purposes, even with the cellphone being used as a router.

VI. Wired Network Experiments

The wired portion of our IoT sensor network consists of the components on the right side of Figure 1, that is the Raspberry Pi with its Sensor Agent, the cloud-hosted web service that receives sensor data, and the cloud-hosted sensor database that stores the sensor data. We ran experiments on four different configurations having different security and scalability attributes.

A. Configuration 1: Base System, Minimal Security, with no Scalability

In this configuration, an Amazon Web Services (AWS) EC2 running Ubuntu server is used for hosting the MySQL sensor database and the web service running on a Glassfish server, as proposed in [5]. The Sensor Agent communicates with the web service through unsecured HTTP messages.

We judge confidentiality, integrity, availability and scalability to be poor in this configuration (Table II). There is no confidentiality or data integrity because the HTTP connection is unsecured and could be monitored or spoofed. Availability is limited because the single AWS instance and its internet endpoint provide a single point of failure. Some awkward limited vertical scaling is possible since the instance can be shut down and replaced by one with more compute power. But the service and database cannot be smoothly horizontally scaled by adding more instances.

TABLE I
TIME TAKEN IN MILLISECONDS FOR THE DATA TRANSMISSION PER MESSAGE

Trial	Medium: Home Router (delay per message in ms)			Medium: Cellphone Hotspot (delay per message in ms)		
	No Encryption	AES CBC Single Encryption	AES CBC Double Encryption	No Encryption	AES CBC Single Encryption	AES CBC Double Encryption
1	0.169	0.2	0.336	0.231	0.253	0.363

2	0.183	0.208	0.314	0.238	0.265	0.361
3	0.169	0.202	0.331	0.252	0.277	0.357
4	0.168	0.189	0.317	0.238	0.281	0.366
5	0.18	0.223	0.31	0.242	0.269	0.368

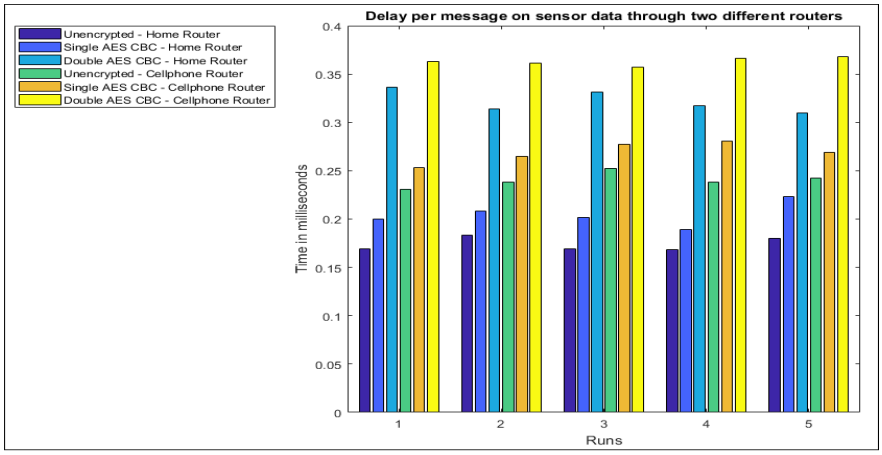

Fig. 4. Comparison of latencies for wireless configurations

B. Configuration 2: Message Security, with No Scalability

In this configuration, we host the web service and the sensor database as before, on a single AWS instance using the same operating system, application server and database server. However, the Glassfish application server is now configured for HTTPS (SSL) messaging utilizing a certificate signed by a well-known certification authority.

In this configuration moderate confidentiality and data integrity are provided by the encrypted SSL messaging. The signed certificate guarantees that messages are being sent to the correct endpoint. However, there is still no authentication of the Sensor Agent, so the web service has no such guarantee as to the source of the data it is receiving.

Availability and scalability continue to be poor, for the same reasons given for configuration 1.

C. Configuration 3: Message Security and Agent Authentication, with No Scalability

This configuration builds on the previous one by assigning credentials to the Sensor Agent consisting of a username and password. These are sent as part of each message and are used by the web service to authenticate the source of incoming messages. Authentication is done using the same MySQL database server which is given an additional table to hold usernames and hashed passwords.

We would rate the confidentiality and data integrity characteristics of this configuration as good (Table II). It shares with configuration 2 the encrypted SSL messaging. Additionally, authentication is now bidirectional, with both certificate authentication of web service to Sensor Agent and password authentication of Sensor Agent to web service. Thus, the web service now has some assurance of the integrity of the data it receives. If, for example, a Sensor Agent were compromised, its credentials could be revoked to limit the damage it could do.

The availability and scalability characteristics of this configuration continue to be poor, due to the continued reliance on a single hardware instance.

D. Configuration 4: Commercial Security and Scalability

Several companies now provide computing infrastructure that allow developers to create software applications that are both secure and scalable. Our fourth configuration relies on such an infrastructure, the collection of Amazon Web Services (AWS).

In this configuration the Sensor Agent, instead of transmitting to the web service directly, sends its messages to an HTTPS endpoint for AWS Kinesis Streams. This is a service for real time handling of large streams of data records [30]. Capacity can be scaled up or down on-the-fly by adding or removing "shards", each providing a certain volume of message processing capacity [31].

TABLE II
SUMMARY OF INFORMATION SECURITY ATTRIBUTES

Config.	Confidentiality	Integrity	Availability	Scalability
#1	poor	poor	poor	poor
#2	moderate	moderate	poor	poor
#3	good	good	poor	poor
#4	good	good	good	good

The web service is replaced by a Java function, compiled into a jar file and deployed to the AWS Lambda service [32]. AWS Lambda is an autoscaling serverless compute service that runs event-based code fragments.

Lambda polls our Kinesis stream and triggers a call to the Java function when a batch of records is available for processing.

When triggered, the Java function reads input records from the Kinesis stream and writes them to a table in AWS DynamoDB, a managed NoSQL database service which acts as the sensor database. DynamoDB provides availability by automatically distributing traffic across multiple servers and by automatically replicating data across independent "availability zones" thus avoiding a single point of failure [33]. As with Kinesis, read/write capacity of a DynamoDB table can be scaled up or down on-the-fly.

Underlying Kinesis, Lambda and DynamoDB is the AWS Identity and Access Management (IAM) service. This is a service that controls access to AWS resources by providing authentication and authorization functions [34]. The Sensor Agent authenticates to Kinesis using IAM-issued credentials. Similarly, the Lambda function uses an IAM-issued role to read from our Kinesis stream and to write to our DynamoDB table.

We rate the confidentiality and data integrity of configuration 4 as good since, combined with the use of HTTPS for messaging, the IAM-based authentications guarantee data confidentiality and integrity all the way from the Sensor Agent to the sensor database. Similarly, we rate availability as good, since all the AWS services claim high availability based on redundancy in their underlying infrastructure. (For highest availability, AWS encourages users to distribute their applications across multiple AWS availability zones, whose data centers are independent of each other [35]. We did not explore this kind of redundancy in our experiments.) Scalability is also good, since each service incorporates automatic or semi-automatic on-the-fly scaling.

E. Experimental Procedure

For the wired network experiments, simulated sensor data messages were generated at the Sensor Agent on the Raspberry Pi Hub and transmitted to the web service. All four configurations were tested.

Each experimental run lasted approximately 300 seconds. The following parameters were varied giving 4 runs for each configuration that is 16 runs in total:

- Data rate: *low rate*, 5 messages per second or *max rate*, as fast as the Raspberry Pi could send
- Packet size: 24 bytes or 1024 bytes.

The clocks on the Raspberry Pi and the servers involved in the experiments were synchronized using NTP services, which allows for measured time differences to be attributable to processing and transmission latencies and not clock differences. The following times in milliseconds were collected for each message:

t1 - message sent from Hub
t2 - message received at the web service
t3 - write to data store completed

The primary measure for evaluating each configuration is end-to-end latency, defined as t3 - t1. This measure would probably be the chief concern of an IoT system designer since it is the time between data generation and availability for processing. Secondary measures include transmission latency, defined as t2 - t1, and processing latency, defined as t3 - t2. These allow us to identify the relative importance of delays incurred in transmitting to the cloud as compared to delays occurring during processing within the web service and data store.

F. Experimental Results

Tables III and IV show the results of the 16 experimental runs. Comparing configurations 1 and 2 we can see that the average end-to-end latency penalty for using SSL message security is quite modest. For example, using the maximum data rate and the 1024-byte messages the average time goes from 19.2 ms to 23.3 ms. Similarly, comparing configuration 2 and 3 the penalty for agent authentication remains small. Using the same example of maximum data rate and 1024 byte records we go from 23.3 ms to 29.9 ms.

The situation changes radically when we compare configuration 4, with the commercial security and stability, to the earlier configurations. The end-to-end latency goes from 29.9 ms to values of over 5 seconds.

TABLE III
AVERAGE END-TO-END LATENCY (MS)

Data Rate	Low Rate		Max. Rate	
Message Size (bytes)	24	1024	24	1024
Config. 1 - Base, Minimal Security, No Scalability	20.8	20.6	18.8	19.2

Config. 2 - Message Security, No Scalability	20.8	23	18.4	23.3
Config. 3 - Message Security, Agent Authentication, No Scalability	21.8	20.2	25.6	29.9
Config. 4 - Commercial Security and Scalability	2101	1696	12425.6	5899

TABLE IV
AVERAGE PROCESSING LATENCY (MS)

Data Rate	Low Rate		Max. Rate	
Message Size (bytes)	24	1024	24	1024
Config. 1 - Base, Minimal Security, No Scalability	0.11	0.09	0.06	0.05
Config. 2 - Message Security, No Scalability	0.06	0.07	0.04	0.06
Config. 3 - Message Security, Agent Authentication, No Scalability	0.23	0.63	0.16	0.25
Config. 4 - Commercial Security and Scalability	295	130	2118	1357

To identify the source of the difference we can look at the two components of end-to-end latency, that is, the transmission latency and the processing latency. Table IV shows the average processing latency on the web service. Except for configuration 4, processing latency is almost zero, indicating that processing time is essentially trivial. This shows that the end-to-end delays come almost entirely in transmitting data from the Sensor Agent to the web service.

For configuration 4, while it is still true that most of the end-to-end latency is transmission latency, the processing latency is more significant, averaging between 130 ms to more than two seconds depending on the data rate and message size.

For many IoT applications the variability of latency may be just as important as the average latency. Variability is most easily seen graphically. Figure 5 plots the end-to-end latency in the maximum data rate runs.

The picture is broadly the same across configurations 1, 2, and 3. The first message in each run has a large latency, presumably owing to the cost of establishing the network connection. This initial latency is higher when HTTPS is used in configurations 2 and 3, which is expected because the session encryption key must be negotiated before transmission begins.

After the initial peak, each plot shows a roughly constant latency but with spikes from time to time. These presumably represent moments when resources are deallocated and reallocated in the underlying computing infrastructure as connections are being re-established.

For configuration 4, with the commercial infrastructure, the picture is very different. The plot shows broad increases and decreases over the 300-second period of the run. Also, we see a spiky, almost sawtooth-like behavior. Finally, there is substantial difference between processing of the 24-byte and the 1024-byte messages.

The greater complexity of the behavior of configuration 4 is perhaps natural considering the greater complexity of the processing needed in a commercial IoT infrastructure. Consider just one aspect, authorization. In configuration 3, authorization is accomplished simply by including a username and password in the Sensor Agent's message to the web service. The web service sends a query to a database table located on the same node and access is authorized if the hashed passwords match.

In the commercial AWS infrastructure, authorization happens multiple times for each message. The Sensor Agent presents credentials to Kinesis to request the right to use the stream. The Lambda function presents credentials to Kinesis to request the right to pull data from the stream, and to DynamoDB to request permission to write to the data store. In each case the Identity and Access Management (IAM) service is consulted to create a context around each request, to consult the different policy documents that govern authorization, and to arrive at a conclusion as to whether each request should be allowed or denied [36]. Compared to configuration 3, the IAM authorization process involves a lot of overhead and several opportunities for variability.

Variability is not due only to authorization and IAM. Kinesis, Lambda, and DynamoDB are also likely to require multiple

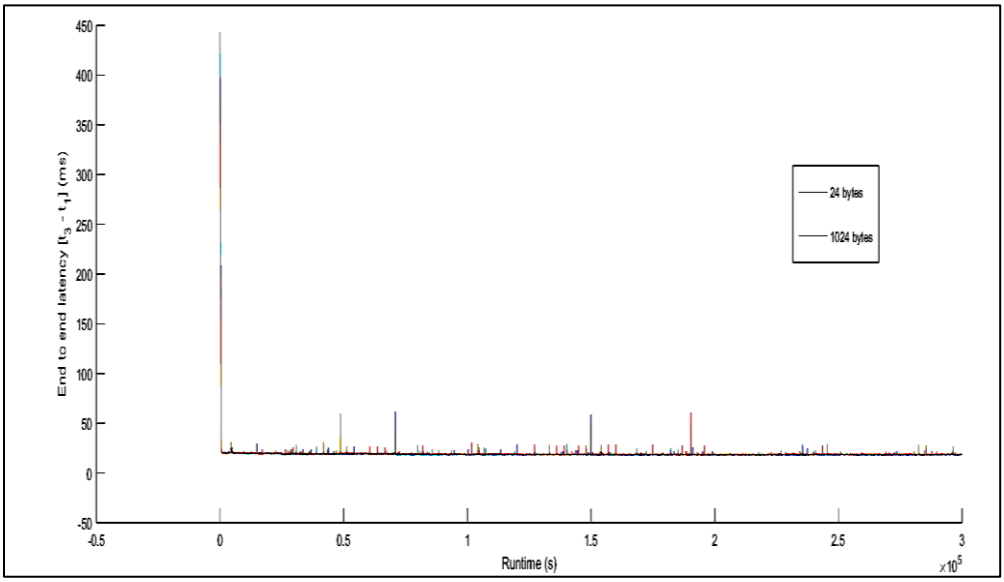

IoT Wired Network Configuration 1

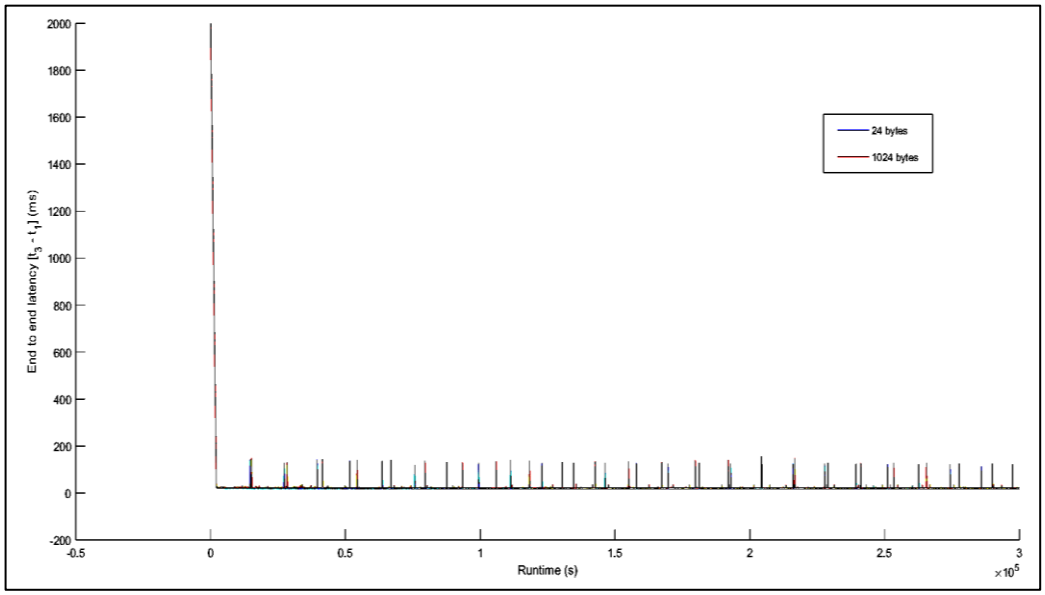

IoT Wired Network Configuration 2

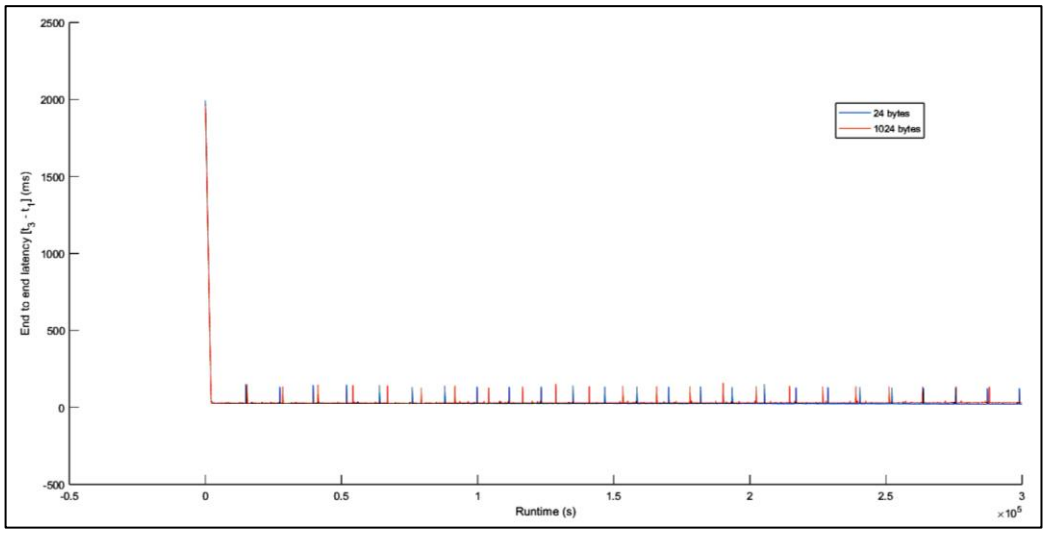

IoT Wired Network Configuration 3

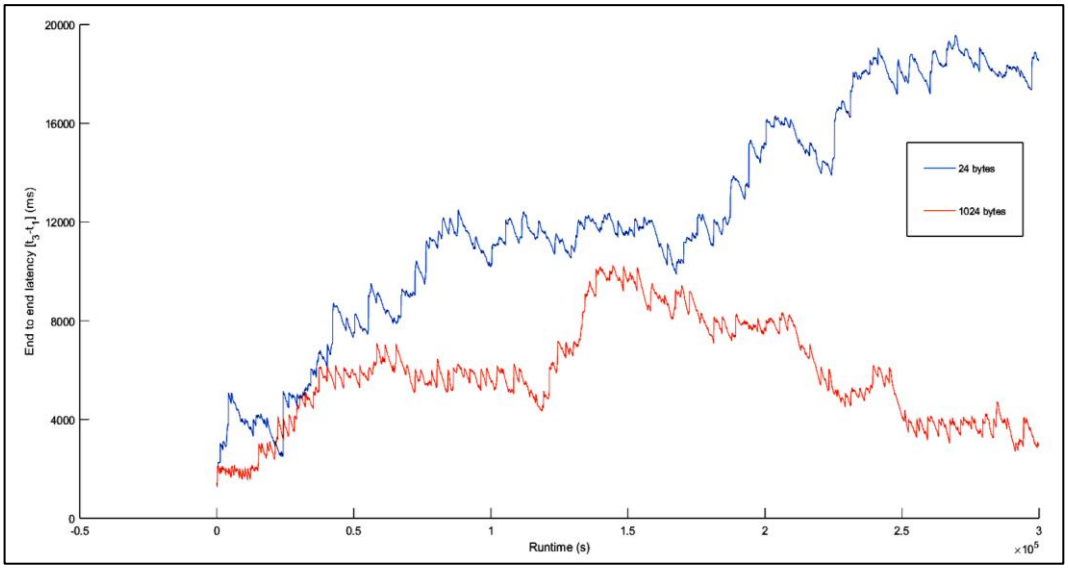

IoT Wired Network Configuration 4

Fig. 5. End-to-End Latency Plots - Max Data Rate
(Note: Plots have Different Vertical Scales)

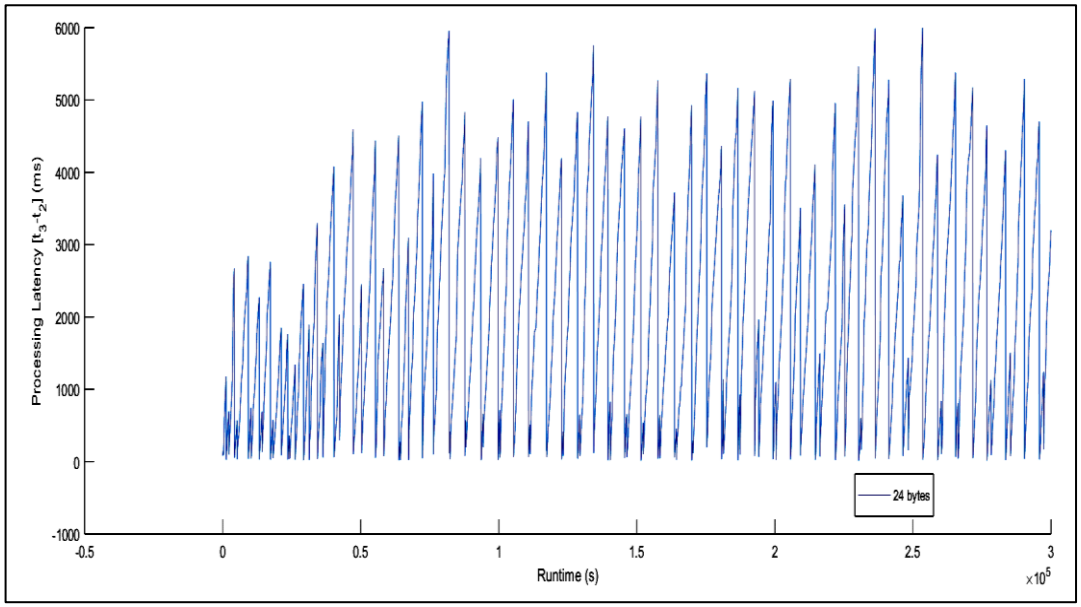

Fig. 6. Configuration 4 - Processing Time for 24-byte records

message communications between multiple nodes to fulfil each task. These nodes are shared across the AWS user base, so processing time can be affected by concurrent actions of other AWS users sharing the same infrastructure.

The results suggest that AWS services likely use buffering and batching strategies to improve average response time, but with the consequence that latency times can have additional sawtooth-like variability as shown in Figure 6. This figure shows the time that configuration 4's web service required to process each sensor message, from message receipt until the DynamoDB service reported that the corresponding data message had been written. As can be seen, the processing latency varied systematically from near zero to 6 seconds. This kind of variability has also been reported by Bermbach in his benchmark studies of cloud services and should perhaps be taken as normal by IoT application developers [20].

VII. Discussion & Conclusions

Security implications create obstacles for wider acceptance of Smart Homes. As the Smart Home market grows, the attack surface for an IoT network within the home grows with it. Smart Home applications face attacks ranging from snooping on transmissions and traffic analysis or leak of message contents, to alteration, fabrication, disruption of communications through node-capture,

routing attacks, or flooding. In several commercially available devices, sensitive data pertaining to parameters from the human body are transmitted over wireless or wired links with little or no security. The goal of this study was to explore the trade-offs in security solutions for IoT applications including Smart Home systems.

Our cloud-based model measured processing overhead while improving on security in the prototype network and adding more scalable services to the system. Our results hinge heavily on the effectiveness of the employed tools as well as the Amazon AWS services. Alternative implementations may yield different performance findings on the efficacy of communicating sensor data to the Web Store with further processing. However, outside the productivity limitations of the code and hardware we used, the current tools and frameworks used in our experiments are industry standard and widely used. Thus, we hope that our results will be representative of what might be encountered in other Smart Home or IoT systems with regards to the price of security and scalability.

Real-time constraints on data collected in a Smart Home or any other IoT Application will vary based on the specifics of the application. The aim of this study was to assess the impact of security overhead on real-time availability of data. The experiment has shown a dramatic difference between custom-built solutions and the use of scalable, secured commercial infrastructure solutions such as Amazon AWS provides. The overhead of adding security in customized solutions as illustrated in configuration 1, 2, and 3 was minimal compared to Amazon AWS as measured in configuration 4. There was also a great variability in latency when using Amazon AWS.

In developing real-time systems, it is vital to consider the trade-off between scalability and robust security versus performance in adopting commercial solutions. As our experiments have shown the greater variability can make a difference in how real-time IoT systems respond to events occurring in smart homes and other IoT applications. In most cases security cannot be compromised but the platform can be chosen in accordance with the real-time constraints of the application.

REFERENCES

[1] A. Pantelopoulos and N. G. Bourbakis. A survey on wearable sensor-based systems for health monitoring and prognosis. IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), 40(1):1-12, 2010.

[2] D. L. Hoffman and T. P. Novak, Thomas, “How to Market the Smart Home: Focus on Emergent Experience, Not Use Cases, Social Science Research Network, January 2016. Available at SSRN: http://ssrn.com/abstract=2840976.

[3] A. Bosche, D. Crawford, D. Jackson, M. Schallehn and C. Schorling “Unlocking Opportunities in the Internet of Things”, Bain and COmpany, January 2016. Available at:
https://www.bain.com/contentassets/5aa3a678438846289af59f62e62a3456/bain_brief_unlocking_opportunities_in_the_internet_of_things.pdf/.

[4] I. Ennajjar, Y. Tabii, and A. Benkaddour, "Security in cloud computing approaches and solutions", 2014 Third IEEE International Colloquium in Information Science and Technology (CIST), Tetouan, Morocco, 20-22 Oct. 2014, DOI: 10.1109/CIST.2014.7016594

[5] T. Reichherzer, A. Mishra, E. Kalaimannan, and N. Wilde, ‘A Case Study on the Trade-Offs between Security, Scalability, and Efficiency in Smart Home Sensor Networks', IEEE: 2016 International Conference on Computational Science and Computational Intelligence (CSCI), 20 March 2017

[6] T. Reichherzer, S. Satterfield, J. Belitsos, J. Chudzynski, L. Watson, “An Agent-Based Architecture for Sensor Data Collection and Reasoning in Smart Home Environments for Independent Living,” In Khoury, R., Drummond, C. (Eds) Lecture Notes in Computer Science, Vol. 9673, Springer-Verlag, pp. 15 – 20, 2016.

[7] NIST, Guide for Conducting Risk Assessments, SP 800-30 Rev. 1, September 2012.

[8] Z. Lu, Y. Zhou, "The Evaluation Model for Network Security", 2014 Fourth International Conference on Communication Systems and Network Technologies, Bhopal, India, 7-9 April 2014, DOI: 10.1109/CSNT.2014.145

[9] M. Botterman, “Internet of Things: An Early Reality of the Future Internet”, Report of the Internet of Things Workshop, for the European Commission, Information Society and Media Directorate General, Networked Enterprise & RFID Unit – D4, Prague, Czech Republic, May 2009.

[10] R. H. Weber, “Internet of Things – New security and privacy challenges”, Elsevier computer law & security review 26, pp. 23 - 30, 2010.

[11] D. Giusto, A. Iera, G. Morabito, and L. Atzori (Eds.), “The Internet of Things”, Springer, 2010. ISBN: 978-1-4419-1673-0.

[12] National Intelligence Council, “Disruptive Civil Technologies – Six Technologies with Potential Impacts on US Interests Out to 2025”, Conference Report CR 2008-07, April 2008.

[13] David Harries and Peter M. Yellowlees. Telemedicine and e-Health. January 2013, 19(1): 61-66. doi:10.1089/tmj.2012.0022.
Replaces:
http://hitconsultant.net/2016/01/28/hackers-caused-98-of-healthcare-data-breaches/.

[14] M. Niland, “Virus Disrupts Train Signals”, The Associated Press, 2003, Available here:
http://www.cbsnews.com/stories/2003/08/21/tech/main569418.shtml.

[15] A. Greenberg, “Hackers Remotely Kill a Jeep On the Highway - With Me In It”, Wired Magazine, July, 2015, Available here: http://www.wired.com/2015/07/hackers-remotely-kill-jeep-highway/.

[16] B. Krebs, “FBI: Smart Meter Hacks Likely to Spread”, Krebs on Security, 2012, Available here: http://krebsonsecurity.com/2012/04/fbi-smart-meter-hacks-likely-to-spread/

[17] Pablo Valerio, "Your Baby Monitor Just Collapsed the Internet", The ICT Scoop, Oct 25, 2016, Available here: https://theictscoop.com/your-baby-monitor-just-collapsed-the-internet-655f1d034dd9

[18] D. Chen and H. Zhao, “Data security and privacy protection issues in cloud computing,” in Proceeding of the International Conference on Computer Science and Electronics Engineering (ICCSEE '12), vol. 1, pp. 647–651, Hangzhou, China, March 2012.

[19] Q. Jing, A.V. Vasilakos, J. Wan, J. Lu and D. Qiu, “Security of the Internet of Things: perspectives and challenges”, Wireless Networks, Vol. 20, Issue 1, pp. 2481-2501, 2014.

[20] D. Bermbach, Quality of Cloud Services: Expect the Unexpected, IEEE Internet Computing, Vol. 21, No. 1, pp. 68-72, Jan. 2017, doi={10.1109/MIC.2017.1}

[21] S. Müller, D. Bermbach, S. Tai and F. Pallas, “Benchmarking the performance impact of transport layer security in cloud database systems”, IEEE International conference on Cloud Engineering (IC2E), Boston, MA, 2014.

[22] S. Kardas¸, S. C¸ elik, M. A. Bingol, and A. Levi, “A new security ¨ and privacy framework for RFID in cloud computing,” in Proceedings of the 5th IEEE International Conference on Cloud Computing Technology and Science (CloudCom '13), Bristol , UK, 2013.

[23] D. Sun, G. Chang, L. Sun, and X. Wang, “Surveying and analyzing security, privacy and trust issues in cloud computing environments,” Procedia Engineering, Vol. 15, No. 1, pp. 2852–2856, 2011.

[24] R. Zhang, O. Berder, J. Gorce, and O. Sentieys, “Energy–delay tradeoff in wireless multihop networks with unreliable links”, Ad hoc Networks, Vol. 10, No. 1, pp. 1306-1321, 2012.

[25] A. Tasiopoulos, C. Tsiaras, and S. Toumpis, “Optimal and achievable cost/delay tradeoffs in delay-tolerant networks”, Computer Networks, Vol. 70, No. 1, pp. 59-74, 2014.
[26] B. Mukherjee, S. Wang, W. Lu, R. Neupane, D. Dunn, Y. Ren, Q. Su, and P. Calyam, “Flexible IoT security middleware for end-to-end cloud-fog communication, Future Generation Computer Systems, Vol. 87, No. 1, pp. 688-703, 2018.
[27] A. Tewari, and B.B. Gupta, “Security, privacy and trust of different layers in Internet-of-Things (IoTs) framework, Future Generation Computer Systems”, Article in Press, 2018.
[28] A. Mishra and D. P. Agrawal, "Energy Conservation and Lifetime Optimization of Wireless Body Sensor Networks for 24x7 Physiological parameters’ Monitoring," Journal of Communications, vol. 10, no. 9, pp. 685-695, 2015. Doi: 10.12720/jcm.10.9.685-695
[29] K. Grgić, I. Špeh, I. Heđi, "A web-based IoT solution for monitoring data using MQTT protocol", 2016 International Conference on Smart Systems and Technologies (SST), Osijek, Croatia, 12-14 Oct. 2016, DOI: 10.1109/SST.2016.7765668
[30] Amazon, "What Is Amazon Kinesis Streams?", http://docs.aws.amazon.com/streams/latest/dev/introduction.html, Link accessed August 20, 2017.
[31] Amazon, Amazon Kinesis Streams Key Concepts, http://docs.aws.amazon.com/streams/latest/dev/key-concepts.html. Link accessed August 28, 2017.
[32] Amazon, "AWS Lambda - Product Details", https://aws.amazon.com/lambda/details/, Link accessed August 20, 2017.
[33] Amazon, Amazon DynamoDB Developer Guide API Version 2012-08-10, http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/dynamodb-dg.pdf, Link accessed August 24, 2017
[34] Amazon, AWS Identity and Access Management User Guide, 2017, http://docs.aws.amazon.com/IAM/latest/UserGuide/iam-ug.pdf. Link accessed August 24, 2017
[35] Amazon, Global Infrastructure, https://aws.amazon.com/about-aws/global-infrastructure/. Link accessed August 28, 2017.
[36] Amazon, Understanding How IAM Works, https://docs.aws.amazon.com/IAM/latest/UserGuide/intro-structure.html. Link accessed May 31, 2018.